

TITLE OF THE INVENTION

METHOD AND DEVICE FOR PERFORMING A QUERY ON A MARKUP DOCUMENT TO CONSERVE MEMORY AND TIME

CROSS REFERENCE TO RELATED APPLICATIONS

[001] This application is based on and hereby claims priority to European Application No. 00125159.4 filed on November 17, 2000 in Europe, the contents of which are hereby incorporated by reference.

BACKGROUND OF THE INVENTION

[002] The invention relates to a method for performing a query on a document created using a Markup language and to software and hardware configured to carry out the method. More specifically, the invention enables the time required to perform a query to be reduced and enables the size of the memory required to perform the query to be reduced as compared to the related art.

[003] There are two basic ways to interface a parser with an application, namely, using an object-based interface and an event-based interface. A Markup language that is becoming popular at the time of writing this application is XML (Extensible Markup Language), and two types of interfaces have been developed for use with XML. The DOM (Document Object Model) interface is an object-based interface and the SAX (Simple Application Programming Interface) is an event-based interface. Related art methods of searching a Markup document using either of these interfaces involve constructing a tree representing the document to be searched.

[004] With a parser using an object-based interface, such as the DOM, the parser explicitly builds a tree of objects that contains all of the elements of the XML document. In contrast, a SAX parser usually accepts a document handler that receives callbacks invoked by the SAX parser. The callbacks inform the document handler of events that are read by the SAX parser. Such events can be, for example, a start-tag and an end-tag. The sequence of callbacks allows the document handler to build a tree of objects of all of the XML elements as they appear in the XML document. However, constructing such a tree requires a great deal of memory and time, and a query, typically, runs several times over the constructed tree.

SUMMARY OF THE INVENTION

[005] It is accordingly an object of the invention to provide a method and a device which overcomes the hereinafore-mentioned disadvantages of the heretofore-known methods and devices of this general type in such a way that the time required to perform a query of a markup document (a document containing data and markup) can be reduced and the size of the memory required to perform the query can be reduced.

[006] With the foregoing and other objects in view there is provided, in accordance with one aspect of the invention a method of performing a query on a Markup document, which includes steps of receiving a query and designing a plurality of filters to reflect a structural linkage of a condition tree representing the query. The step of designing the plurality of filters includes designing a highest-level filter that can become active only if an event-based parser indicates that an element for which the highest-level filter is searching has been found. The step of designing the plurality of filters also includes designing a lowest-level filter that can become active only when the highest-level filter has become active and when the parser indicates that an element for which the lowest-level filter is searching has been parsed. The method also includes a step of parsing a Markup document, and a step of checking the lowest-level filter to determine whether it has found the element for which it has been searching.

[007] A query is expressed as a condition tree, which has at every single condition a linkage to a filter, as described above. A single condition determines its result by evaluating its linked filter. A composite condition determines its value by evaluating all of its sub-conditions.

[008] In accordance with an added feature of the invention, the step of designing the plurality of filters includes: designing at least one intermediate-level filter that can become active only when the highest-level filter has become active and when the parser indicates that an element for which the intermediate-level filter is searching has been parsed; and designing the lowest-level filter to become active only when the intermediate-level filter has become active.

[009] In accordance with an additional feature of the invention, the lowest-level filter is defined as a first lowest-level filter; and the method includes steps of designing a second lowest-level filter that can become active only when the highest-level filter has become active and when the parser indicates that an element for which the lowest-level filter is searching has been parsed;

and checking the second lowest-level filter to determine whether it has found the element for which it has been searching.

[0010] In accordance with another feature of the invention, the value filter is designed to become active only when the highest-level filter has become active and when the parser indicates that an element for which the value filter is searching has been parsed. If the first lowest-level filter has found the element for which it has been searching and the second lowest-level filter has found the element for which it has been searching, an element is obtained from the value filter that is linked to the elements in the first lowest-level filter and in the second lowest-level filter.

[0011] In accordance with a further feature of the invention, the method includes designing a value filter that will become active only when the highest-level filter has become active and when the parser indicates that an element for which the value filter is searching has been parsed; and if the lowest-level filter has found the element for which it has been searching, obtaining an element from the value filter that is linked to the element in the lowest-level filter.

[0012] In accordance with a further added feature of the invention, computer executable instructions for performing the method are stored on a computer-readable medium.

[0013] In accordance with a concomitant feature of the invention, a computer device is programmed to perform the method by executing the instructions that have been stored on a computer readable medium.

[0014] One aspect of the invention enables desired information to be read from a Markup document in an extremely efficient manner and involves using an event-based interface to read the document such that a tree need not be constructed representing the Markup document.

BRIEF DESCRIPTION OF THE DRAWINGS

[0015] These and other objects and advantages of the present invention will become more apparent and more readily appreciated from the following description of the preferred embodiments, taken in conjunction with the accompanying drawings of which:

Fig. 1 shows an XML document named package.csd;

Fig. 2 shows the interaction of methods necessary to perform the simple XQL query-"softpkg/implementation/@id" on the XML document; and

Fig. 3 is a diagram showing the hierarchy of the filters used to perform the complex XQL query-"softpkg/implementation @id.

DESCRIPTION OF THE PREFERRED EMBODIMENTS

[0016] Reference will now be made in detail to the preferred embodiments of the present invention, examples of which are illustrated in the accompanying drawings, wherein like reference numerals refer to like elements throughout.

[0017] One aspect of the invention involves using an event-based interface to read a Markup document. An exemplary embodiment of the invention will be described that uses a SAX interface to read an XML document. However, it should be apparent that the invention could be constructed using another event-based interface constructed for use with another Markup language, and therefore, the invention should not be construed as being limited to use with XML documents.

[0018] One aspect of the invention is based upon the concept of constructing a condition tree representing the query to be performed on the document and constructing filters in accordance with the tree, instead of constructing a tree of document elements beforehand. The filters are document handlers that are hierarchically registered with each other and at the topmost level with a parser. The filter cascade begins with the construction of forwarding filters to narrow the elements to read from. A query filter is created which also serves as a forwarding filter. During the creation of the query filter, the condition expression, as part of the query, is read and a condition cascade is initialized. The condition cascade uses the composite design pattern to represent the conditions and their Boolean links. After construction of the query filter, a filter chain for a value filter is created. At the bottom level, this is mostly, an "existence", "elementlist", or "attribute" filter which serves as a value filter. If a query filter was created due to the presence of a condition in the query, this value filter is linked to the condition. If the query did not contain a condition, the value filter serves as the filter from which the results are directly obtained. The topmost filter is registered with the parser, for example, an XML parser supporting the SAX interface. The topmost filter then delegates to all of the lower level filters, including the "query", "existence", "elementlist", "forwarding" and/or "attribute" filters. The query

filter evaluates the condition at certain check points which are at the end of its designated scope. In the example provided below, the evaluation would be at the end of the element "implementation". If there are composite conditions, they would evaluate their sub-conditions based upon the Boolean expressions that link them together. Finally, if the condition is evaluated to be true, the associated value filter is read.

[0019] Fig. 1 shows an example of an XML document or descriptor named package.csd that is used to describe CORBA components. Fig. 2 shows the interaction of methods or operations necessary to perform a simple query of the XML document. The XQL (XML Query Language) statement: "softpkg/implementation/@id" - can be used to query the document for the "id" attribute of an "implementation" element that is a child of a "softpkg" element. The query can be represented as a tree showing the sequence of events from "softpkg" to "implementation" and finally to "id". Specifically, "id" is a child of "implementation" which is a child of "softpkg".

[0020] Referring to Fig. 2, one will see the creation of methods used to implement the filters performing the query. The filters are registered hierarchically. The forwarding filter "softpkg" 2 is registered with an SAX parser 4 and will be activated upon receiving a callback from the parser 4 indicating that a "softpkg" event has been read. The forwarding filter "implementation" 6 is registered with the forwarding filter "softpkg" 2 such that the filter "implementation" 6 can receive callbacks from the parser 4 only after the filter "softpkg" 2 has been activated. The filter "implementation" 6 will be activated upon receiving a callback indicating that an "implementation" event has been activated. The attribute filter "AttributeFilter" 8 is registered with the filter "implementation" 6 such that the filter "AttributeFilter" 8 can receive callbacks from the parser 4 only after the filter "implementation" 6 has been activated. The filters 2, 6, 8 are, in effect, SAX document handlers.

[0021] After the filters 2, 6, 8 have been created and properly registered, the document to be queried, in this example, "package.csd" is parsed. After parsing the document, the "getlength" method is performed to see if the "AttributeFilter" 8 has obtained one or more results in response to the query, and if so, the "getResult" method is performed to obtain one or more results from the "AttributeFilter" 8.

[0022] Because the filters 2, 6, 8 receive callbacks from the parser 4 in response to the events as they are being read by the parser 4, and because the filters 2, 6, 8 are hierarchically

registered, the filters 2, 6, 8 enable a query to be performed on the document without having to produce a tree representing the document. In effect, a query tree is continually applied to the elements of the document as the document is being parsed. The filters 2, 6, 8 act to "filter out" the event or events that are of interest in response to the query, if in fact, at least one such event exists in the document. It can be seen that the condition expression only has to be "parsed" once for queries to any number of different XML Markup documents.

[0023] An example of a complex query will now be discussed. Referring to Fig. 3, one will see the hierarchy of filters that can be used to perform the complex XQL query-
"softpkg/implementation @id" on the XML document. The forwarding filter "softpkg" 10 is registered with the SAX-parser and will become active only when a "softpkg" element is found. The forwarding filter "implementation" 12 is registered with the filter "softpkg" 10 and can become active only when the filter "softpkg" 10 is active and when an "implementation" element is found. A first hierarchical filter chain 14 is registered with the filter "implementation" to find "os" elements having name attributes of 'WinNT' where these "os" elements are also children of "implementation" elements. A second hierarchical filter chain 16 is registered with the filter "implementation" 12 to find "compiler" elements having name attributes of 'MSVC' where these "compiler" elements are also children of "implementation" elements. The leftmost filter shown in Fig. 3 is an attribute filter that is used as a value filter 18 to temporarily store "id" attributes of "implementation" elements. The "name" attribute filters are checked to see if the desired elements have been found. If the "name" attribute filter in the first filter chain 14 has found an element for which it is searching, and if the "name" attribute filter in the second filter chain 16 has found an element for which it is searching, then the necessary composite condition is satisfied and the one or more "id" attributes in the value filter 18 are obtained from the value filter 18 in response to the query.

[0024] The computer language C++, for example, could be used to construct computer executable instructions that would implement the filters, and the computer executable instructions could be stored on a computer readable medium, such a ROM (read only memory) or a RAM (random access memory). The computer executable instructions could also be stored on a portable computer disk for downloading into a computer device at a later time, wherein the computer device, upon executing the instructions, would perform the method described hereinabove.

[0025] The invention has been described in detail with particular reference to preferred embodiments thereof and examples, but it will be understood that variations and modifications can be effected within the spirit and scope of the invention.